

10/20/00
jc950 U.S. PTO

10/23/00

A

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

ATTY DOCKET NO.: 5577-222
DATE: October 20, 2000

jc929 U.S. PTO
09/693268
10/20/00

**UTILITY PATENT APPLICATION TRANSMITTAL LETTER
AND FEE TRANSMITTAL FORM (37 CFR 1.53(b))**

BOX PATENT APPLICATION
Assistant Commissioner for Patents
Washington, DC 20231

Sir:

Transmitted herewith for filing under 37 CFR 1.53(b) is:

- ☒ a patent application
☐ a Continuation ☐ a Divisional ☐ a Continuation-in-Part (CIP)
of prior application no.: ; filed
☐ A Small Entity Statement(s) was filed in the prior application; Status still proper and desired.

Inventor(s) or Application Identifier: **Mike E. Baskey; Roy F. Brabson; Lap T. Huynh; Peter Yocum**

Entitled: **METHODS, SYSTEMS AND COMPUTER PROGRAM PRODUCTS
FOR SERVER BASED TYPE OF SERVICE CLASSIFICATION OF
A COMMUNICATION REQUEST**

Enclosed are:

1. ☒ Application Transmittal Letter and Fee Transmittal Form (*A duplicate is enclosed for fee processing*)
2. ☒ 37 pages of Specification (including 43 claims)
3. ☒ Five (5) sheets of Formal Drawings (35 USC 113)
4. ☐ Oath or Declaration
 - a. ☐ newly executed (*original or copy*)
 - b. ☐ copy from prior application (37 CFR 1.63(d) (*for continuation/divisional*) [Note Box 5 Below]
 - c. ☐ DELETION OF INVENTOR(S) (*Signed statement deleting inventor(s) named in the prior application*)
5. ☐ Incorporation By Reference (*useable if box 4b is checked*)

The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
6. ☐ Microfiche Computer Program (*Appendix*)
7. ☐ Assignment papers (*cover sheet(s) and document(s)*)
8. ☐ Small Entity Statement(s)
9. ☒ Information Disclosure Statement, PTO-1449, and 4 references cited
10. ☐ Preliminary Amendment (*Please enter all claim amendments prior to calculating the filing fee.*)
11. ☐ English Translation Document
12. ☐ Certified Copy of Application No. ; Filed

10/20/00



- ☐ Sequence Listing/ Sequence Listing Diskette
- a. ☐ computer readable copy
- b. ☐ paper copy
- c. ☐ statement in support
- ☐ An Associate Power of Attorney
- ☒ Return Receipt Postcard (MPEP 503) (Should be specifically itemized)
- ☐ Other:

The fee has been calculated as shown below:

	Column 1 No. Filed	Column 2 No. Extra	Small Entity Rate Fee	Large Entity Rate Fee
BASIC FEE			\$355.00	\$710.00
TOTAL CLAIMS	43- 20 =	23	x 09 = \$	x 18 = \$414.00
INDEP CLAIMS	4- 3 =	1	x 40 = \$	x 80 = \$80.00
<input type="checkbox"/> MULTIPLE Dependent Claims Presented			+ 135 = \$	+ 270 = \$
If the difference in Col. 1 is less than zero, Enter "0" in Col. 2			Total \$	Total: \$1204.00

- ☐ A check in the amount of \$ _____ to cover the filing fee is enclosed.
- ☒ Please charge my Deposit Account No. 09-0461 in the amount of \$1,244.00 to cover the filing fee (this amount includes \$40.00 for Recordation of Assignment).
- ☒ The Commissioner is hereby authorized to charge payment of the following fees associated with this communication or credit any overpayment to Deposit Account No. 09-0461.
- ☐ Any additional filing fees required under 37 CFR 1.16.
- ☐ Any patent application processing fees under 37 CFR 1.17.

Respectfully submitted,

Robert W. Glatz
Registration No. 36,811

Correspondence Address:
USPTO Customer Number: **20792**
Myers Bigel Sibley & Sajovec, P.A.
Post Office Box 37428
Raleigh, NC 27627
Tel (919) 854-1400
Fax (919) 854-1401

CERTIFICATE OF EXPRESS MAILING

Express Mail Label No. EL682672608US
Date of Deposit: October 20, 2000

I hereby certify that this correspondence is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to Box Patent Application, Assistant Commissioner For Patents, Washington, DC 20231.

Michele P. McMahan
Date of Signature: October 20, 2000


ATTORNEY DOCKET NO.: 5577-222

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Filed: Concurrently Herewith
For: METHODS, SYSTEMS AND COMPUTER PROGRAM PRODUCTS FOR
SERVER BASED TYPE OF SERVICE CLASSIFICATION OF A
COMMUNICATION REQUEST

Date: October 20, 2000



Box Patent Application
Commissioner for Patents
Washington, DC 20231

APPLICATION FILED UNDER 37 CFR 1.41(c)

Sir:

The above identified application is being filed on behalf of the inventor(s), Mike E. Baskey, resident of Wappingers Falls, NY, Roy F. Brabson, resident of Raleigh, NC, Lap Thiet Huynh, resident of Apex, NC, Peter Yocum, resident of Wappingers Falls, NY, under the provisions of 37 CFR 1.41(c). A Declaration and Power of Attorney from the inventor will follow, 37 CFR 1.63.

Respectfully submitted,

A handwritten signature in black ink, appearing to read "Robert W. Glatz".

Robert W. Glatz
Registration No. 36,811

Myers Bigel Sibley & Sajovec
P.O. Box 37428
Raleigh, NC 27627
Tel (919) 854-1400
Fax (919) 854-1401

CERTIFICATE OF EXPRESS MAILING

"Express Mail" mailing label number: EL 682672608US
Date of Deposit: October 20, 2000

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to Box Patent Application, Commissioner For Patents, Washington, DC 20231.

A handwritten signature in black ink, appearing to read "Michele P. McMahan".
Michele P. McMahan Date of Signature: October 20, 2000

METHODS, SYSTEMS AND COMPUTER PROGRAM PRODUCTS FOR
SERVER BASED TYPE OF SERVICE CLASSIFICATION OF A
COMMUNICATION REQUEST

Related Applications

The present application is a continuation-in-part
of Application Serial No. 09/645,651, filed August 24,
2000, entitled "Methods, Systems and Computer Program
Products for Processing an Event Based on Policy Rules
5 Using Hashing" (Attorney Docket No. 5577-197) which is
incorporated herein by reference as if set forth in its
entirety.

Field of the Invention

10 The present invention relates to type of service
classification and more particularly to type of service
classification associated with communications over a
network.

Background of the Invention

15 Computer networks have grown increasingly complex
with the use of distributed client/server applications,
mixed platforms and multiple protocols all in a single
physical backbone. The control of traffic on networks
20 is likewise moving from centralized information systems

Such an approach may fail to integrate all the desired elements to achieve consistent response time, for example, when processing web-based transactions which benefit from prioritization to achieve consistent response times. For example, such web-based requests are all typically managed at the same priority level which may result in downloads, browses and business transactions being managed at the same priority level with both network and server resources being applied equally across what may preferably be treated as three different priorities of workload. Such a result may be encountered in conventional systems as all such web-based requests typically travel through the network and arrive at the server using the same port (generally port 80 or port 443 if secured socket layer (SSL) communications are being used).

An example of such a QoS product is the Web Traffic Express product available from IBM corporation of Armonk, New York which generally provides web associated quality of service in a separate device positioned between the server/client (endpoint) device and the communication network. This separate device may then determine network qualities of service but generally does not synchronize its setting with the actual server/client device providing an application using web communications and/or database server functions. Thus, the Web Traffic Express product generally does not provide full correlation between network and server resources. Similarly, web QoS products available from Hewlett Packard Corporation and Cisco Corporation generally do not integrate with server-based workload management processes. They also typically execute on the application level.

Summary of the Invention

Embodiments of the present invention include methods, systems and computer program products for type
5 of service classification of a communication request for an application executing on a server. An application plug-in associated with the application is provided in an operating system kernel of the server. The application plug-in receives the communication
10 request and obtains application level information from the received communication request. The application plug-in further assigns a type of service classification to the received communication request based on the obtained application level information and
15 provides the assigned type of service classification information for the communication request to a process executing on the server for processing communications from the server responsive to the communication request.

20 The communication request may be a TCP/IP protocol communication in which case the application level information from the received communication request is level 5 or above information. The assigned type of service classification information for the
25 communication request may then be provided to a TCP/IP kernel executing on the server.

In further embodiments of the present invention, the information provided to the TCP/IP kernel includes one or more of a source IP address, a destination IP
30 address and a TCP/IP port number associated with the communication request as an associated identification for the assigned type of service. The TCP/IP kernel in such embodiments associates the assigned type of service with a plurality of communications from the
35 server responsive to the received communication request

based on the provided source IP address, destination IP address and TCP/IP port number. For example, an HTTP request may request transfer of a "page" which is large enough to extend across multiple outgoing communication
5 packets from the server. Each of these outgoing communication packets may provide the designated type of service determined from the incoming communication request.

In particular embodiments, the assigned type of
10 service classification may be associated with one or more new thread instances initiated on the server based on the obtained application level information. In further embodiments, the assigned type of service classification may be associated across sockets API for
15 associated connections based on the obtained application level information.

In other embodiments of the present invention, the type of service classification is assigned based on workload management information associated with the
20 server. Information associated with the received communication may be provided to a workload management process executing on the server and the workload management information may then be received from the workload management process. The assigned type of
25 service classification may assign one or more of a central processing unit (CPU) priority allocation, a memory allocation or an input/output (I/O) bandwidth allocation to the received communication request.

In further embodiments of the present invention,
30 type of service information is included in communications from the server responsive to the communication request based on the assigned type of service classification. The type of service information is usable by a network communicating the
35 communications from the server for prioritization of

traffic flows on the network. The type of service information may be selected to provide substantially uniform priority in handling the communication request by both the server and the network or may provide a different type of service for network prioritization of communications from the server than the assigned type of service classification provides from the server for processing of the communication request. For example, a separate connection may be provided for communications from the server responsive to the communication request to support the different type of service for network prioritization.

The communication request may be a web-based request and the application may be a web application. The web-based request may be a hypertext transport protocol (HTTP) request and the application level information may include a universal resource locator (URL) of the HTTP request. The URL in such embodiments may then be parsed on the server to make a policy based determination of the type of service classification for the communication request. The application level information may include an identification of a user initiating the communication request at a source device of the communication request. In further embodiments, the application level information may be obtained from a cookie contained in the communication request. For example, the identification of a user initiating the communication request may be obtained from a cookie in the communication request.

In other embodiments of the present invention, systems are provided for type of service classification of a communication request for an application executing on a server. The system includes a communication process executing on the server that processes communications between the server and a communication

network associated with the communication request based on an associated type of service classification. The system further includes an application plug-in process associated with the application in an operating system kernel of the server that obtains application level information from the received communication request, assigns the type of service classification to the received communication request based on the application level information and provides the assigned type of service classification to the communication process. In particular embodiments, the system further includes a workload management process executing on the server that receives information about the communication request from the application plug-in process and provides information related to server resources for use in allocating server resources to the communication process for use in processing communications between the server and a communication network associated with the communication request based on an associated type of service classification.

While the invention has been described above primarily with respect to the method aspects of the invention, both systems and/or computer program products are also provided.

Brief Description of the Drawings

Figure 1 is a block diagram of a network environment in which the present invention may be implemented;

Figure 2 is a block diagram of data processing systems according to embodiments of the present invention;

Figure 3 is a more detailed block diagram of data processing systems according to embodiments of the present invention;

Figure 4 is a flowchart illustrating operations according to embodiments of the present invention; and

Figure 5 is a flowchart illustrating operations according to further embodiments the present invention.

Detailed Description of Preferred Embodiments

10 The present invention now will be described more fully hereinafter with reference to the accompanying drawings, in which preferred embodiments of the invention are shown. This invention may, however, be embodied in many different forms and should not be
15 construed as limited to the embodiments set forth herein; rather, these embodiments are provided so that this disclosure will be thorough and complete, and will fully convey the scope of the invention to those skilled in the art.

20 As will be appreciated by one of skill in the art, the present invention may be embodied as a method, data processing system, or computer program product. Accordingly, the present invention may take the form of an entirely hardware embodiment, an entirely software
25 embodiment or an embodiment combining software and hardware aspects. Furthermore, the present invention may take the form of a computer program product on a computer-usable storage medium having computer-usable program code means embodied in the medium. Any suitable
30 computer readable medium may be utilized including hard disks, CD-ROMs, optical storage devices, a transmission media such as those supporting the Internet or an intranet, or magnetic storage devices.

Computer program code for carrying out operations of the present invention may be written in an object oriented programming language such as Java®, Smalltalk or C++. However, the computer program code for carrying out operations of the present invention may also be written in conventional procedural programming languages, such as the "C" programming language. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, on partly on the user's computer and partly on a remote computer. In the latter scenario, the remote computer may be connected to the user's computer through a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

The present invention is described below with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions specified in the flowchart and/or block diagram block or blocks.

These computer program instructions may also be stored in a computer-readable memory that can direct a computer or other programmable data processing apparatus to function in a particular manner, such that
5 the instructions stored in the computer-readable memory produce an article of manufacture including instruction means which implement the function specified in the flowchart and/or block diagram block or blocks.

The computer program instructions may also be
10 loaded onto a computer or other programmable data processing apparatus to cause a series of operational steps to be performed on the computer or other programmable apparatus to produce a computer implemented process such that the instructions which
15 execute on the computer or other programmable apparatus provide steps for implementing the functions specified in the flowchart and/or block diagram block or blocks.

Referring first to the block diagram illustration of **Figure 1**, a network environment in which the present
20 invention may be implemented will be generally described. As illustrated in **Figure 1**, the communication network **100** includes a plurality of connecting nodes and endpoint nodes. As shown in **Figure 1**, two endpoint servers **105** and two clients **145**
25 are shown which are endpoints on the communication network **100**. However, additional devices may be connected and a single computer device may serve as both a server and a client in different transactions and may further function as a connecting node between
30 network **100** and another network. Accordingly, as used herein, the term "server" will refer to an endpoint node with respect to a communication request as contrasted with a router or bridge device.

As shown in **Figure 1**, endpoint servers **105** connect through a bridge switch **115** and a router **120** to the IP network **125**. Clients **145** are also connected to the IP network **125** through a bridge/switch **115** and a router

5 **120**. Additional bridge/switches **115** and routers **120** may be included in the IP network **125** as will be understood by one of ordinary skill in the art.

The communication network **100** may be provided by wired or wireless links using a variety of physical

10 layer media and still stay within the teachings and scope of the present invention. Furthermore, while the present invention is described primarily in the context of communication networks and control of communication traffic on such networks, the present invention is not

15 so limited and may be utilized beneficially in processing events based on type of service classifications, particularly in environments where processing speed for automatically identifying and executing an appropriate type of service is important.

20 For example, fast methodologies for classifying Web based communication requests (such as those designated by an HTTP address) for assignment of appropriate QoS based on policy rules is desirable for use with communication servers handling communications over

25 networks such as the communication network **100**. The routers **120** may also operate to provide appropriate service levels to packets routed through the respective routers **120** based on network policy rules.

Type of service classification for processing of

30 communication requests may be beneficially implemented in endpoint devices such as servers **105** or clients **145**. Such capabilities may be particularly beneficial in light of emerging network trends where service differentiation of different data types is important in

00000T" 89266960

guaranteeing service level agreement (SLA) performance. Various embodiments of the present invention, as will be described herein, may provide communication servers which classify a Web request based, for example, on a policy definition, which classification may be used to assign an appropriate type of service for a request once a matching policy rule is found. This approach may reduce the overall number of policy rules that need to be evaluated for each event and processing of rules efficiently and in real time may be implemented for processing Web requests.

The present invention will now be further described with reference to the block diagram of **Figure 2** which illustrates data processing systems according to embodiments of the present invention. As illustrated in **Figure 2**, the system **230** may include input device(s) **232** such as a keyboard or keypad, a display **234**, and a memory **236** that communicate with a processor **238**. The memory **236** can include, but is not limited to, the following types of devices: cache, ROM, PROM, EPROM, EEPROM, flash memory, SRAM, and DRAM. The data processing system **230** may further include a storage system **242**, a speaker **244** and an I/O data port(s) **246** that also communicate with the processor **238**. The storage system **242** may include removable and/or fixed media such as floppy disks, ZIP drives, hard disks or the like as well as virtual storage such as a RAMDISK. The I/O data port(s) **246** can be used to transfer information between the data processing system **230** and another computer system or a network (e.g., the Internet). Such data processing systems may include, for example, personal computers, laptop computers, mainframe computers, pervasive computing devices such as personal digital assistants, smartphones or the

like, or even embedded processing systems. The components of a particular data processing system may be conventional or custom components, such as those used in many conventional computing devices, which may
5 be configured to operate as described herein.

Figure 3 is a block diagram of a data processing system that illustrates systems, methods, and computer program products in accordance with embodiments of the present invention. As shown in the embodiments of
10 **Figure 3**, the server (or client device acting as a server in that it is receiving communication requests which result in outgoing communications initiated by an application executing thereon) **305** includes application level processes and processes executing in the
15 operating system kernel. As shown in **Figure 3**, two exemplary applications supporting communications over the communication network **325**, identified as application **335** and application **340**, are illustrated. Furthermore, a policy agent **345** and workload management process **350** are shown as application level processes on
20 the server **305**. For the illustration of **Figure 3**, the communicating applications **335**, **340** are shown as communicating with the communication network **325** using a TCP/IP protocol. Thus, a TCP/IP kernel **355** is shown
25 as being included in the operating system kernel of the server **305**. It is to be understood that other kernel elements **375** are typically also found on the server **305**, which other kernel elements **375** may further support communications with applications over the
30 communication network **325** using other layered protocols.

The TCP/IP kernel **355**, as shown in **Figure 3**, communicates with the communication network **325** through

a router/bridge device **320**. Where the router/bridge device **320** is a router, it typically routes communications packets based on level 4 information, such as destination IP address. Where the
5 router/bridge device **320** is a bridge, it may process communications within a network node using level 3 information, such as machine identification number, rather than level 4 IP address information.

As shown in **Figure 3**, the TCP/IP kernel **355**
10 includes a communication process **360**. The communication process **360** executes on the server **305** to process communications between the applications **335**, **340** of the server **305** and the communication network **325**. Such communications are typically initiated
15 responsive to an incoming communication directed to one of the applications **335**, **340** which communication initiates subsequent outgoing communications from the server **305** based on content provided by the respective applications **335**, **340** responsive to the incoming
20 communication request.

Thus, the communication process **360** may support incoming communication request processing as well as processing of responsive communications from the server **305**. More particularly, in accordance with embodiments
25 of the present invention, the communication process **360** processes such outgoing communications from the server **305** based on an associated type of service classification defined for such communications.

An application plug-in process **365** is provided
30 that assigns the type of service classification. A single application plug-in process **365** is illustrated in **Figure 3** which is associated with the application **335** as indicated by the dotted line of **Figure 3**. It is

to be understood that in various embodiments of the present invention, a separate application plug-in process may be provided for each of a plurality of applications supporting communications with the communication network **325**. Alternatively, an application plug-in process **365** may be associated with a type of application, a plurality of instances of which type of application may be executing on the server **305**.

10 The application plug-in process **365** obtains application level information from the received communication request. For example, in the context of a TCP/IP protocol communication as illustrated in **Figure 3**, the application plug-in process **365** obtains
15 level 5 or above information from the received communication request. The application plug-in process **365** further assigns the type of service classification to the received communication request based upon the application level information it obtains and provides
20 the assigned type of service classification to the communication process **360**.

Also shown in the TCP/IP kernel **355** is a fast cache **370**, such as a fast response cache accelerator (FRCA), which may support communications, such as web
25 type communications. Such caches are known, for example, for use with HTTP type requests based upon a URL associated with a "page" request. The URL information is an example of application level information which may be utilized by the application
30 plug-in process **365** to assign a type of service. Accordingly, operations of the application plug-in process **365** may be beneficially coordinated with those associated with the fast cache **370** in processing cache supported communications.

from Apple Computer, UNIX or Linux, proprietary operating systems or dedicated operating systems, for example, for embedded data processing systems.

Operations according to various embodiments of the present invention will now be described further with reference to the flowchart illustrations of **Figures 4-5**. Operations begin with reference to **Figure 4** at block **400** when an application plug-in associated with an application executing on a server is provided in an operating system kernel of the server. The application plug-in subsequently receives communication requests including a communication request for the associated application executing on the server (block **405**). Application level information from the received communication request is obtained by the application plug-in (block **410**). For example, with reference to a TCP/IP protocol communication, the application level information would correspond to level 5 or above information obtained from the received communication request.

A type of service classification is assigned to the received communication request based on the obtained application level information (block **415**). The assigned type of service classification information for the communication request is provided to a process executing on the server for processing communications from the server responsive to the communication request (block **420**). For example, the assigned type of service classification information for the communication request may be provided to a TCP/IP kernel executing on the server which processes communications using the TCP/IP protocol.

In addition to providing the type of service classification information, the application plug-in may

further provide the TCP/IP kernel a source IP address, a destination IP address and a TCP/IP port number associated with the communication request as an identification for the assigned type of service. Using this source and destination IP address information and the port number, the TCP/IP kernel in various embodiments of the present invention associates the assigned type of service with a plurality of communications from the server responsive to the received communication request. Thus, a type of service once assigned to a request, may be applied by the TCP/IP kernel supporting a TCP/IP communication stack to a plurality of output packets initiated responsive to the received communication request.

For example, the TCP/IP kernel acting as a communication process **360** may associate the assigned type of service classification with one or more thread instances initiated on the server based on the obtained application level information. The type of service markings may, thus, be propagated to new threads/processes within the server with those new threads/processes established to include the assigned type of service markings across the sockets API for associated connections.

Where policy based rules are used to assign the type of service classification at block **415** using both source and destination IP address and port number along with application level information (for example, a URL or other classification mechanism), a fine grained differentiation of different service classes by transaction type may be provided in accordance with embodiments of the present invention. In other words, in various embodiments of the present invention, for example, web-based communications embodiments, communication processing transactions may be classified

by parsing the URL on the server to determine a policy-based priority for each class as to how it is to be processed on the server (which may include allocation of resources such as central process unit (CPU) resources, allocation of the memory resources and allocation of input/output (I/O) bandwidth. Furthermore, such prioritization may be maintained across responses sent for each request.

As will now be described with reference to the flowchart illustration of operations according to further embodiments of the present invention in **Figure 5**, such prioritization may also be coordinated with network prioritization by linking the classification policy for network prioritization with the server policies for workload prioritization. Doing so may improve consistency between the ways that required network and server resources are managed so that the most important work may receive preferences both on the server and across the network carrying the communications as soon as the respective prioritized transactions are recognized. In the case of secured transactions, for example, using SSL (port 443), the server prioritization may actually be the first opportunity for a transaction (communication request) to be classified through inspection of the contents of the packets, especially if the transmitting user did not choose to distribute the encryption keys needed to do such content based prioritization to remote platforms. Distribution of such keys to remote platforms could reduce the integrity of the protection provided by the secured communication protocol. Linking the input classification of a communication request and prioritization (for server workload management) to an outbound classification and prioritization (for network resource management) using

the same (or a related) policy may further help provide end-to-end consistent resource management according to the QoS/SLA requirements associated with the communication request.

5 Such operations will now be described with reference to various embodiments in **Figure 5**.

Operations begin at block **500** when application level information is obtained from a received communication request. For the embodiments illustrated in **Figure 5**,
10 information associated with the received communication is provided to a workload management process executing on the server (block **505**). For example, the application plug-in process **365** may provide such information to the workload management process **350**.
15 Workload management information is received from the workload management process responsive to the provided information associated with the received communication (block **510**). The workload manager may thus be utilized on the server to support enforcement of policy rule
20 based type of service classification with respect to resources such as CPU priority, memory and I/O bandwidth allocation consistent with the goals and rights of users associated with the communication request, such as expectations for response time and/or
25 throughput of supported communications. The type of service classification may then be assigned based on the workload management information associated with the server received from the workload management process (block **515**) and provided to the communication process
30 (block **520**).

In a further aspect of various embodiments of the present invention, type of service markings are provided in outbound replies (consistent with the inbound type of service classification generated

responsive to the receipt of a communication request)
to prioritize traffic across the network, such as a
local area network (LAN) and/or a wide area network
(WAN) handling the responsive communications. Such
5 marking of outbound communications, as will be
described, may facilitate balancing resource allocation
consistent with customer policy for workload importance
and coordinating server-based actions with networking
mechanisms like traffic prioritization giving
10 differentiated services based on type of service
markings.

As shown in **Figure 5**, type of service information
is associated with communications from the server which
are generated responsive to the communication request
15 (block **525**). The type of service information is based
on the assigned type of service classification
generated responsive to the received communication
request. As described above, the type of service
information, which is preferably usable by a network
20 communicating the communications from the server for
prioritization of traffic flows on the network, is
incorporated in the responsive communications (block
530). The selection of the type of service information
may be planned to provide a consistent service, such as
25 the same priority, service both at the server and the
network level.

Alternatively, the network type of service
information may be selected to provide a different type
of service for network prioritization of communications
30 from the server than the assigned type of service
classification provides from the server in processing
of the communication request on the server itself. For
example, on the server, outbound type of service
information markings could be overridden by applying a
35 different set of filter rules on the outbound traffic.

By way of further example, an inbound communication request could spawn a separate connection for streaming video as a result of a commercial transaction. More generally, a separate connection may be provided for communications from the server responsive to the communication request to support a type of service for network prioritization of communications from the server which is different from a type of service the server itself provides in processing the communication request.

It is further to be understood that, while the examples above referred to URL as the application level information, other application level information not normally utilized by the communication protocol stack may beneficially be used in keeping with the present invention. For example, an identification of a user initiating the communication request at a source device of the communication request may be determined as the application level information. In a web-type communication context, the application level information could be obtained from a cookie contained in the communication request. In fact, for certain applications, an identification of a user initiating a communication request may itself be contained from a cookie obtained in the communication request.

As described generally above, the various embodiments of the present invention provide for server based quality of service classification. This approach may provide correlation between server and network based prioritization. Furthermore, by providing the application plug-in in an operating system kernel, classification operations may be done efficiently and may further be done in a manner consistent with existing communication caching and classification mechanisms already implemented in a server. The type

of service classification may also be integrated with workload management processes on the server so as to potentially better manage resource and bandwidth allocations at the server.

5 The flowcharts and block diagrams of **Figures 1**
through **5** illustrate the architecture, functionality,
and operation of possible implementations of systems,
methods and computer program products according to
various embodiments of the present invention. In this
10 regard, each block in the flow charts or block diagrams
may represent a module, segment, or portion of code,
which comprises one or more executable instructions for
implementing the specified logical function(s). It
should also be noted that, in some alternative
15 implementations, the functions noted in the blocks may
occur out of the order noted in the figures. For
example, two blocks shown in succession may, in fact,
be executed substantially concurrently, or the blocks
may sometimes be executed in the reverse order,
20 depending upon the functionality involved.

In the drawings and specification, there have been
disclosed typical preferred embodiments of the
invention and, although specific terms are employed,
they are used in a generic and descriptive sense only
25 and not for purposed of limitation, the scope of the
invention being set forth in the following claims.

THAT WHICH IS CLAIMED:

1. A method for type of service classification
of a communication request for an application executing
5 on a server, the method comprising the steps of:

providing an application plug-in associated with
the application in an operating system kernel of the
server;

10 wherein the application plug-in performs the
following steps:

receiving the communication request;

obtaining application level information from the
received communication request;

15 assigning a type of service classification to the
received communication request based on the obtained
application level information; and

providing the assigned type of service
classification information for the communication
request to a process executing on the server for
20 processing communications from the server responsive to
the communication request.

2. The method of Claim 1 wherein the
communication request is a TCP/IP protocol
25 communication and wherein the step of obtaining
application level information from the received
communication request comprises the step of obtaining
level 5 or above information from the received
communication request.

30

3. The method of Claim 2 wherein the step of
providing the assigned type of service classification
information for the communication request to a process
executing on the server comprises the step of providing
35 the assigned type of service classification information

for the communication request to a TCP/IP kernel
executing on the server.

4. The method of Claim 3 wherein the step of
5 providing the assigned type of service classification
information for the communication request to a TCP/IP
kernel further comprises the step of providing a source
IP address, a destination IP address and a TCP/IP port
number associated with the communication request to the
10 TCP/IP kernel as an associated identification for the
assigned type of service and wherein the method further
comprises the step performed by the TCP/IP kernel of
associating the assigned type of service with a
plurality of communications from the server responsive
15 to the received communication request based on the
provided source IP address, destination IP address and
TCP/IP port number.

5. The method of Claim 4 further comprising the
20 step of associating the assigned type of service
classification with a new thread instance initiated on
the server based on the obtained application level
information.

6. The method of Claim 5 wherein the step of
25 associating the assigned type of service classification
with a new thread instance initiated on the server
based on the obtained application level information
further comprises the step of associating the assigned
30 type of service classification across a sockets API for
associated connections based on the obtained
application level information.

7. The method of Claim 2 wherein the step of
35 assigning a type of service classification further

comprises the step of assigning the type of service classification based on workload management information associated with the server.

5 8. The method of Claim 7 further comprising the steps of:

 providing information associated with the received communication to a workload management process executing on the server; and

10 receiving the workload management information from the workload management process.

 9. The method of Claim 7 wherein the assigned type of service classification assigns at least one of
15 a central processing unit (CPU) priority allocation, a memory allocation or an input/output (I/O) bandwidth allocation to the received communication request.

 10. The method of Claim 2 further comprising the
20 step of including type of service information in communications from the server responsive to the communication request based on the assigned type of service classification, the type of service information being usable by a network communicating the
25 communications from the server for prioritization of traffic flows on the network.

 11. The method of Claim 10 wherein the type of service information provides a different type of
30 service for network prioritization of communications from the server than the assigned type of service classification provides from the server for processing of the communication request.

12. The method of Claim 11 wherein the method further comprises the step of providing a separate connection for communications from the server responsive to the communication request to support the different type of service for network prioritization of communications from the server than the assigned type of service classification provides from the server for processing of the communication request.

13. The method of Claim 2 wherein the communication request is a web-based request and the application is a web application.

14. The method of Claim 13 wherein the web-based request is a hypertext transport protocol (HTTP) request and wherein the application level information includes a universal resource locator (URL) of the HTTP request and wherein the assigning step further comprises the step of parsing the URL on the server to make a policy based determination of the type of service classification for the communication request.

15. The method of Claim 14 wherein the assigned type of service classification assigns at least one of a central processing unit (CPU) priority allocation, a memory allocation or an input/output (I/O) bandwidth allocation to the received communication request.

16. The method of Claim 14 wherein the step of providing the assigned type of service classification information for the communication request to a process executing on the server comprises the step of providing the assigned type of service classification information for the communication request to a TCP/IP kernel executing on the server and wherein the method further

received communication request, assigns the type of
service classification to the received communication
request based on the application level information and
provides the assigned type of service classification to
5 the communication process.

21. The system of Claim 20 further comprising a
workload management process executing on the server
that receives information about the communication
10 request from the application plug-in process and
provides information related to server resources for
use in allocating server resources to the communication
process for use in processing communications between
the server and a communication network associated with
15 the communication request based on an associated type
of service classification.

22. A system for type of service classification
of a communication request for an application executing
20 on a server, the system comprising:

a communication process executing on the server
that processes communications from the server
responsive to the communication request; and

an application plug-in associated with the
25 application in an operating system kernel of the
server;

wherein the application plug-in further comprises:
means for receiving the communication request;
means for obtaining application level information
30 from the received communication request;

means for assigning a type of service
classification to the received communication request
based on the obtained application level information;
and

means for providing the assigned type of service classification information for the communication request to the communication process.

5 23. The system of Claim 22 wherein the communication request is a TCP/IP protocol communication and wherein the means for obtaining application level information from the received communication request comprises means for obtaining
10 level 5 or above information from the received communication request.

15 24. The system of Claim 23 wherein the means for providing the assigned type of service classification information for the communication request to the communication process comprises means for providing the assigned type of service classification information for the communication request to a TCP/IP kernel executing on the server.

20 25. The system of Claim 23 wherein the means for assigning a type of service classification further comprises means for assigning the type of service classification based on workload management information
25 associated with the server.

 26. The system of Claim 25 further comprising:
 means for providing information associated with the received communication to a workload management
30 process executing on the server; and
 means for receiving the workload management information from the workload management process.

 27. The system of Claim 23 further comprising
35 means for including type of service information in

communications from the server responsive to the
communication request based on the assigned type of
service classification, the type of service information
being usable by a network communicating the

5 communications from the server for prioritization of
traffic flows on the network.

28. The system of Claim 23 wherein the
communication request is a web-based request and the
10 application is a web application.

29. The system of Claim 28 wherein the web-based
request is a hypertext transport protocol (HTTP)
request and wherein the application level information
15 includes a universal resource locator (URL) of the HTTP
request and wherein the means for assigning further
comprises means for parsing the URL on the server to
make a policy based determination of the type of
service classification for the communication request.

20

30. The system of Claim 29 wherein the assigned
type of service classification assigns at least one of
a central processing unit (CPU) priority allocation, a
memory allocation or an input/output (I/O) bandwidth
25 allocation to the received communication request.

31. The system of Claim 29 wherein the means for
providing the assigned type of service classification
information for the communication request to the
30 communication process comprises means for providing the
assigned type of service classification information for
the communication request to a TCP/IP kernel executing
on the server and wherein the system further comprises
the TCP/IP kernel, wherein the TCP/IP kernel further
35 comprises means for associating the assigned type of

service with a plurality of communications from the server responsive to the received communication request.

5 32. The system of Claim 31 further comprising means for including type of service information in communications from the server responsive to the communication request based on the assigned type of service classification, the type of service information
10 being usable by a network communicating the communications from the server for prioritization of traffic flows on the network.

15 33. A computer program product for type of service classification of a communication request for an application executing on a server, comprising:

 a computer-readable storage medium having computer-readable program code embodied in said medium, said computer-readable program code comprising:
20 application plug-in computer-readable program code associated with the application for execution in an operating system kernel of the server;

 wherein the application plug-in computer-readable program code further comprises:

25 computer-readable program code which receives the communication request;

 computer-readable program code which obtains application level information from the received communication request;

30 computer-readable program code which assigns a type of service classification to the received communication request based on the obtained application level information; and

 computer-readable program code which provides the
35 assigned type of service classification information for

the communication request to a process executing on the server for processing communications from the server responsive to the communication request.

5 34. The computer program product of Claim 33 wherein the communication request is a TCP/IP protocol communication and wherein the computer-readable program code which obtains application level information from the received communication request comprises computer-
10 readable program code which obtains level 5 or above information from the received communication request.

15 35. The computer program product of Claim 34 wherein the computer-readable program code which provides the assigned type of service classification information for the communication request to the communication process comprises computer-readable program code which provides the assigned type of
20 communication request to a TCP/IP kernel executing on the server.

25 36. The computer program product of Claim 34 wherein the computer-readable program code which assigns a type of service classification further comprises computer-readable program code which assigns the type of service classification based on workload management information associated with the server.

30 37. The computer program product of Claim 36 further comprising:

computer-readable program code which provides information associated with the received communication to a workload management process executing on the
35 server; and

computer-readable program code which receives the workload management information from the workload management process.

5 38. The computer program product of Claim 34 further comprising computer-readable program code which includes type of service information in communications from the server responsive to the communication request based on the assigned type of service classification,
10 the type of service information being usable by a network communicating the communications from the server for prioritization of traffic flows on the network.

15 39. The computer program product of Claim 34 wherein the communication request is a web-based request and the application is a web application.

20 40. The computer program product of Claim 39 wherein the web-based request is a hypertext transport protocol (HTTP) request and wherein the application level information includes a universal resource locator (URL) of the HTTP request and wherein the computer-readable program code which assigns further comprises
25 computer-readable program code which parses the URL on the server to make a policy based determination of the type of service classification for the communication request.

30 41. The computer program product of Claim 40 wherein the assigned type of service classification assigns at least one of a central processing unit (CPU) priority allocation, a memory allocation or an input/output (I/O) bandwidth allocation to the received
35 communication request.

METHODS, SYSTEMS AND COMPUTER PROGRAM PRODUCTS FOR
SERVER BASED TYPE OF SERVICE CLASSIFICATION OF A
COMMUNICATION REQUEST

Abstract of the Disclosure

Methods, systems and computer program products are provided for type of service classification of a communication request for an application executing on a server. An application plug-in associated with the application is provided in an operating system kernel of the server. The application plug-in receives the communication request and obtains application level information from the received communication request. The application plug-in further assigns a type of service classification to the received communication request based on the obtained application level information and provides the assigned type of service classification information for the communication request to a process executing on the server for processing communications from the server responsive to the communication request.

000207" 29265960




Figure 1

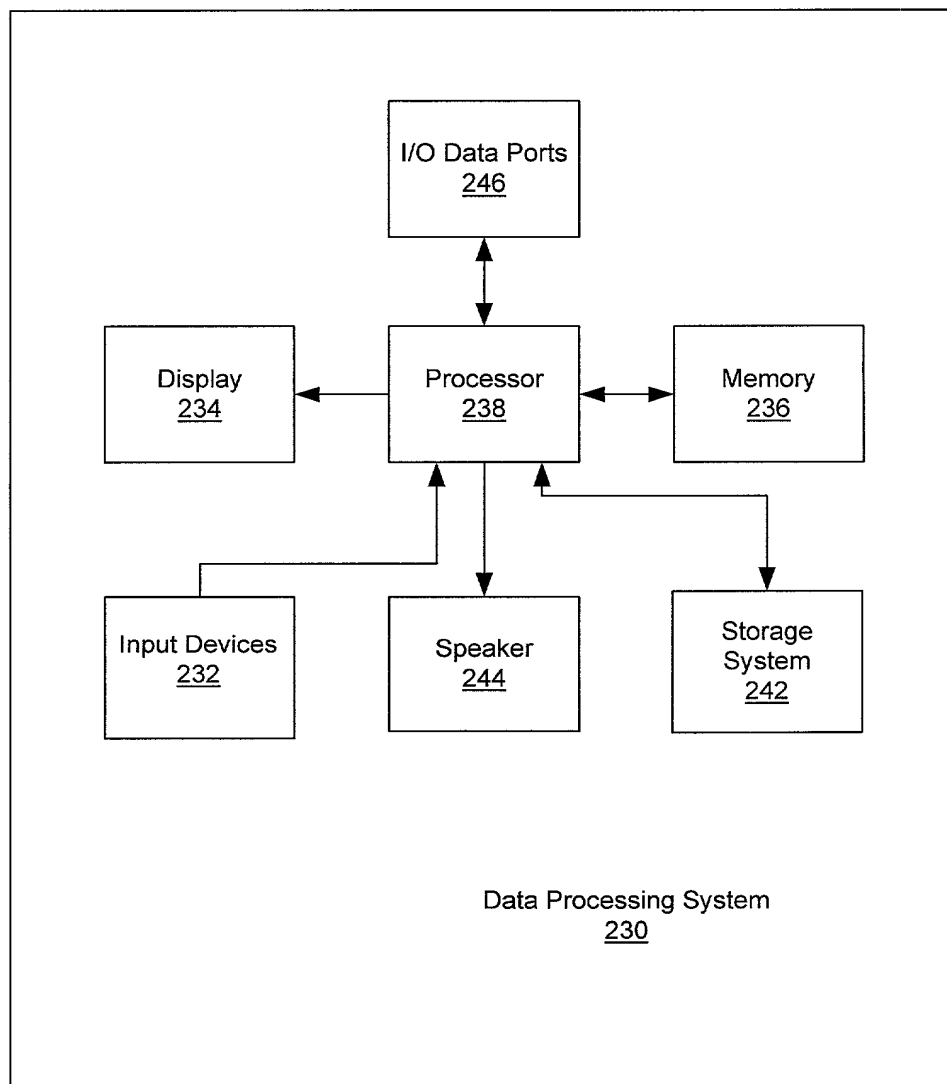


Figure 2

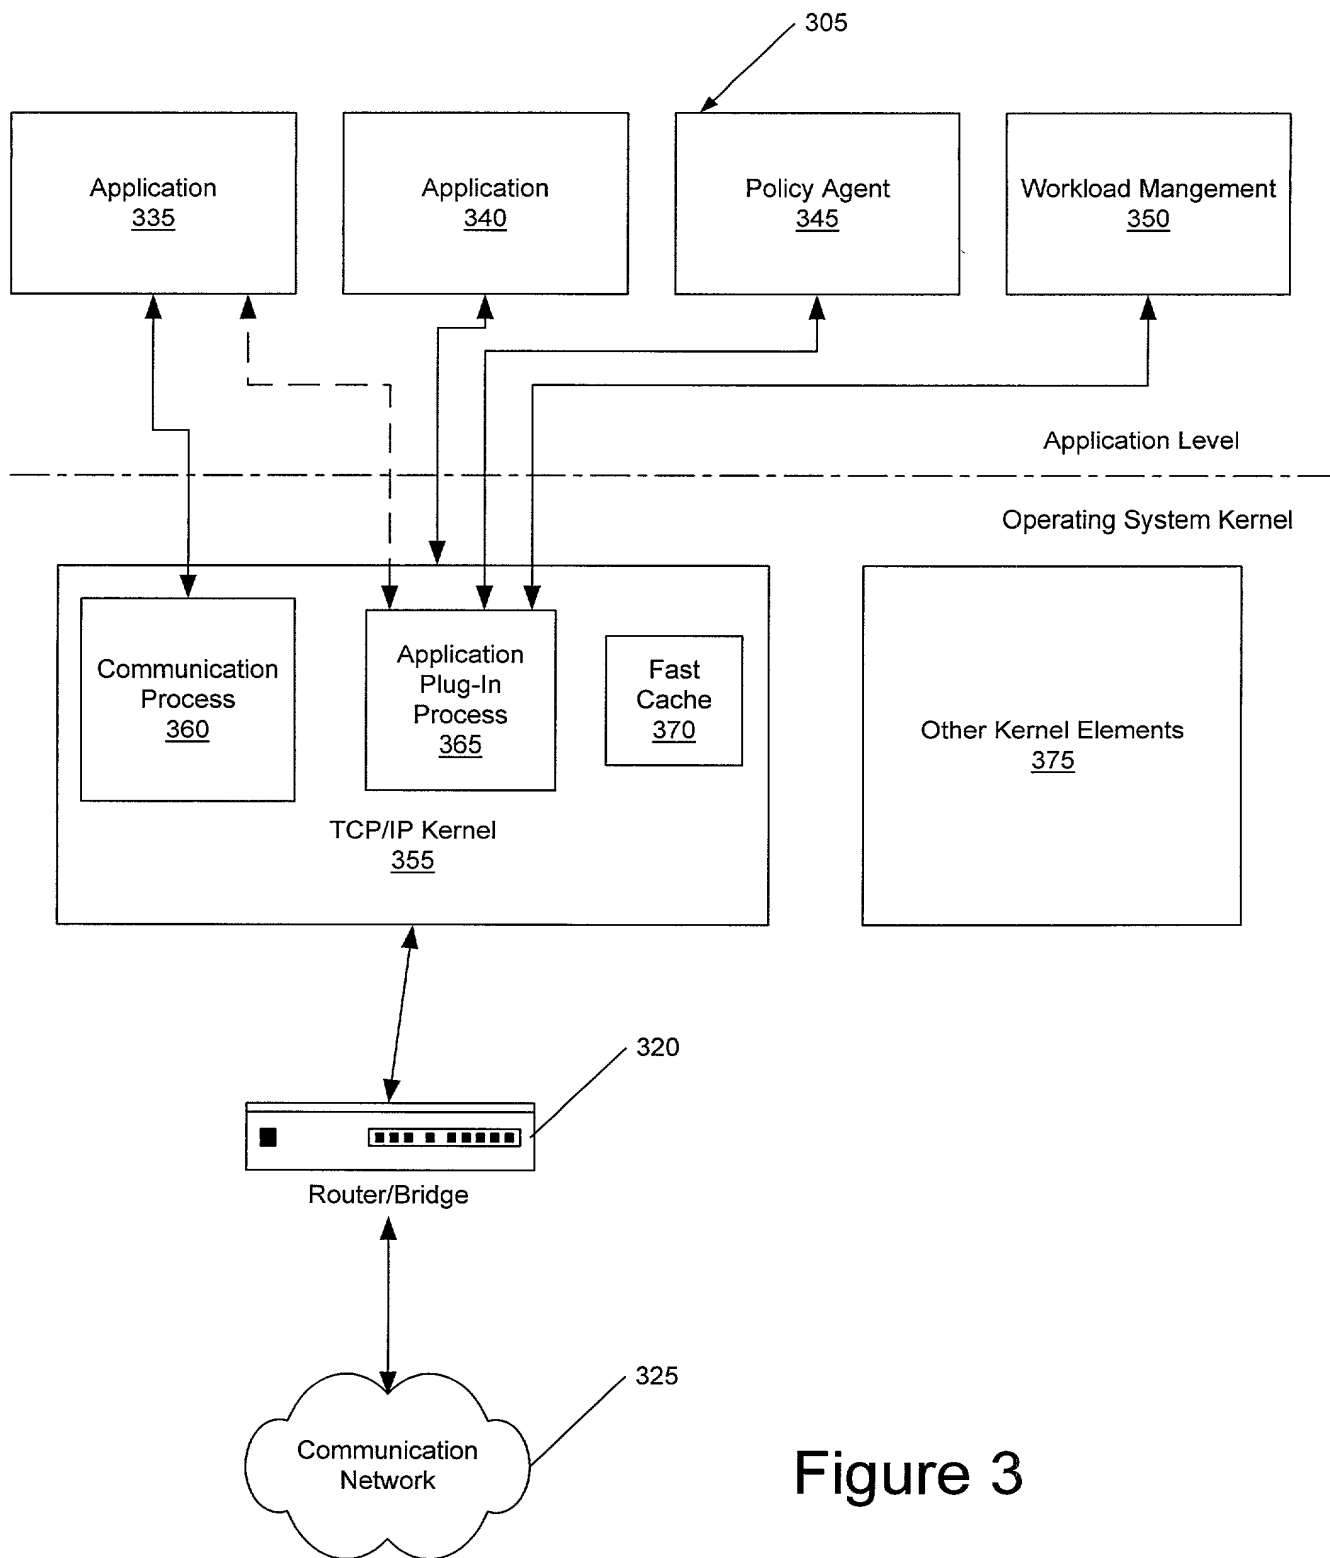


Figure 3

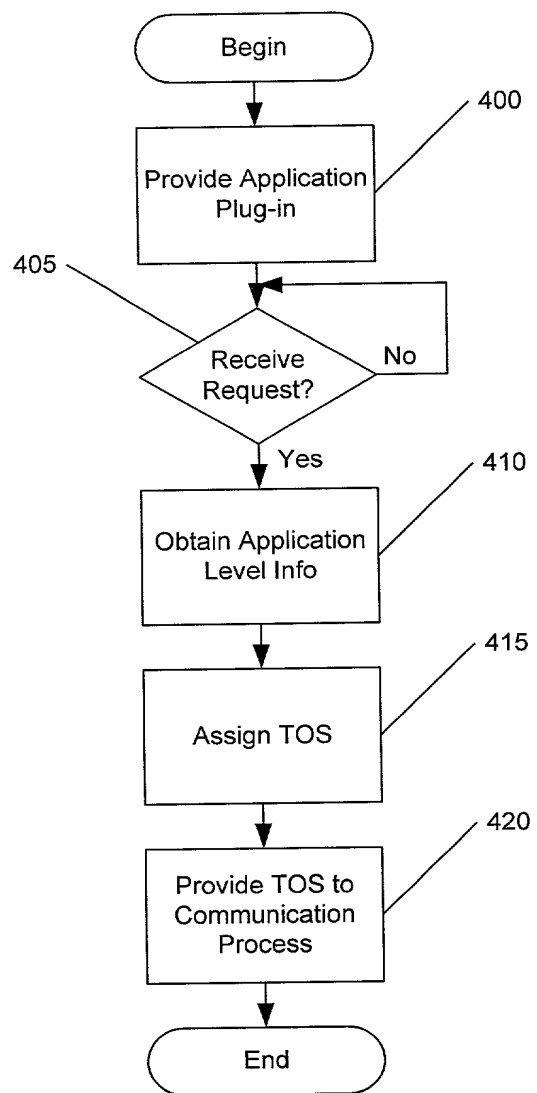


Figure 4

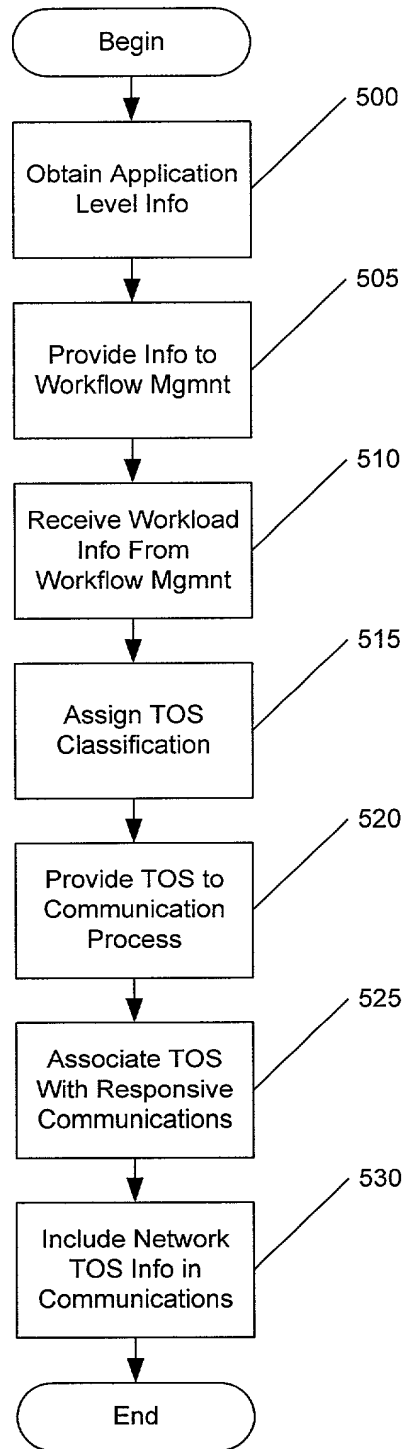


Figure 5